

A METHOD OF CONTROLLING INTERFERENCE FROM A TRANSMITTER IN
ONE COMMUNICATION SYSTEM TO A RECEIVER IN ANOTHER
COMMUNICATION SYSTEM

5 This invention relates to a method of controlling interference from a transmitter in one communication system to a receiver in another communication system. In general, a certain frequency band is allocated for a particular type of communication system, e.g. mobile phones and then within that frequency band, any equipment which wishes to communicate must share the allocated spectrum with other user equipment

10 giving rise to the possibility of interference. In any system with at least one transmitter and one receiver, it is assumed that the transmitter is the cause of the interference. However, the fundamental problem of spectrum sharing is not the transmitters, but the receivers. This is because, with current systems, it is possible to sense the presence or absence of signal from a transmitter that is already operating on a frequency one wishes

15 to share, but it is not possible to determine whether transmission on a particular frequency will cause interference at a receiver or not. The presence of a strong signal might indicate that a remote receiver will have no problem with any interference the transmitting terminal might generate. Equally a weak signal might indicate that a nearby receiver cannot cope with any interference the transmitting terminal might

20 produce. This runs contrary to the standard philosophy of dynamic channel assignment in which the channels with minimum interference are preferred.

In accordance with a first aspect of the present invention, a method of controlling interference from a transmitter in one communication system to a receiver in another communication system comprises transmitting a beacon from a beacon transmitter associated with the receiver representative of a frequency at which the receiver is trying to receive; listening for the beacon at a beacon receiver associated with the transmitter; and deriving a power spectral density limit for a transmission from the transmitter based upon the strength of the beacon received at the beacon receiver.

The problem of whether or not a transmission will cause interference can only be addressed if it is known where the receivers are, so in the present invention, receivers transmit and transmitters receive to convey this information. If a beacon is received above a certain level, then an assessment is made of the extent to which the transmitter power needs to be reduced to have enough power to transmit, without

interfering with the transmitter with which that beacon is associated. Thus, priority can be given to protecting an incumbent system from additional interference, without preventing transmission in the other communication system entirely.

Preferably, for a plurality of beacons received representing the same frequency,
5 the derived maximum transmit power spectral density is related to that of the beacon received at the highest power.

It is assumed that if a transmission is set, such that it does not interfere with the receiver represented by the beacon at the highest power, then all the others will be unaffected too, whether because they are further away or otherwise more tolerant of
10 interference.

Preferably, the method further comprises comparing the derived transmit power spectral density limit with a predetermined minimum transmit power spectral density required by the transmitter for that frequency; and transmitting a signal at that frequency, only if the determined transmit power spectral density limit exceeds the
15 minimum.

If the transmission would not be of sufficient power to reach its intended destination, then it is not sent. This can be overcome by sending via an intermediate terminal, so that the required minimum level is lower.

Preferably, a predetermined maximum transmit power spectral density is set, if
20 no beacons are received at the transmitter.

If the receiver in a sharing system, whose beacon is being received as the strongest beacon, is protected by the sharing transmitter setting its power accordingly, then this guarantees that any other receivers using that frequency experience even less interference.

25 Preferably, the method further comprises choosing a transmission frequency for the transmitter which permits the maximum power spectral density for the transmission.

Preferably, the transmission from the transmitter is transmitted at a frequency derived by determining the strongest received beacon which represents any one
30 frequency; thereafter selecting, from the determined strongest beacons, the beacon with the lowest power; and transmitting at the frequency represented by that selected beacon.

Preferably, a transmit power spectral density for a transmission from the transmitter is set dependent upon the strength of the received beacon at the chosen frequency.

Preferably, the maximum permitted power spectral density of the transmitter is
5 set at the product of the receiver beacon power; and a factor by which the receiver can be de-sensitised minus one; and the resultant of the receiver noise figure divided by the product of the effective bandwidth at the beacon receiver for receiving the beacon, the minimum signal to noise ratio for receiving the beacon in its effective bandwidth and the noise figure of the beacon receiver at the transmitter.

10 Preferably, a random time division multiple access (TDMA) protocol is applied, whereby beacons representing different frequencies transmit at different times, such that over a series of cycles a beacon representing each frequency will be heard at a different time relative to another particular represented frequency, such that no one frequency at a higher power consistently blocks reception of a beacon representing
15 another frequency at a lower power.

Preferably, a code division multiple access (CDMA) protocol is applied, whereby beacons representing different frequencies are distinguished from one another by different codes.

20 Use of a CDMA protocol allows a large number of beacons to share the same frequency band whilst being in band to provide multipath diversity.

Preferably, a correlation period of a CDMA component of the beacon signal is controlled by an FFT controller.

25 Preferably, each beacon transmits a type identifier and each beacon receiver comprises type specific correlation means, such that a beacon receiver can ignore same type beacons in determining whether or not or how much power to transmit.

Preferably, a receiver transmits a beacon only if interference levels exceed an acceptable value.

This saves power in the beacon transmitter, by avoiding unnecessary transmissions.

30 Preferably, the beacon power is adapted to the wanted signal power received at the receiver.

Adaptation of the beacon power allows more interference to be experienced at the beacon transmitting site where this will not disrupt communication.

Preferably, the beacon power is adapted to the interference power received at the receiver.

Adapting the beacon power allows the beacon transmit site equipment to adjust the allowable transmission from multiple sources, such that sum of all interference 5 received is at an acceptable level.

Preferably, a bandwidth managed by a beacon is sufficiently narrow that substantial correlation of shadow fading applies across that bandwidth.

Preferably, beacon communication is separated from a spectrum that is managed by the beacons by using a different frequency.

10 Preferably, each beacon occupies a frequency bandwidth which is small compared with the total bandwidth managed by that beacon.

Preferably, neighbouring beacons in a managed bandwidth manage discrete contiguous sections of frequency; each section comprising a fraction of the beacon managed band, each beacon being separated from the frequency bandwidth which it 15 manages by the alternate fraction.

This arrangement for bandwidth management allows efficient separation of the beacon band from the communications band to allow duplex filtering.

The fractions can be unequal, such as 1/3 and 2/3, but typically each fraction is ½, since this is most efficient fraction to use.

20 Preferably, the beacon receiver is periodically tested with an internal beacon of known power and its associated transmitter is prevented from transmitting if a beacon receiver fault occurs.

Alternatively, beacon and system communications are separated in time.

25 This requires timing control, for example, provided by putting satellite timing receivers in all transmit and receive terminals.

Preferably, beacon reception and transmission happen at the same equipment, separated in time, by arranging for reception to take place whenever transmission is not according to schedules of the random TDMA protocol.

30 This allows equipment that is sharing the spectrum with a beacon generating system to also transmit its own beacons for the benefit of any potential third sharing system.

In accordance with a second aspect of the present invention, a communication system comprises at least one transmitter belonging to one communication system and

at least one receiver belonging to another communication system, wherein a beacon transmitter is associated with the at least one receiver and a beacon receiver is associated with the at least one transmitter, whereby a power spectral density limit for transmission at any one transmitter is determined based upon the strength of the or each beacon received at the associated beacon receiver.

5 In accordance with a third aspect of the present invention, a transmitter for a communication system, the transmitter being provided with an associated beacon receiver, whereby a power spectral density limit for transmission from the transmitter is determined based on the strength of one or more beacons received at the associated 10 beacon receiver.

In accordance with a fourth aspect of the present invention, a receiver for a communication system, the receiver being provided with an associated beacon transmitter, whereby a beacon can be transmitted by the beacon transmitter to control interference, such that a desired maximum power spectral density of interference 15 received at the receiver is achieved.

An example of a method of controlling interference from a transmitter to a receiver in a communication system will now be described with reference to the accompanying drawings in which:

20 Figure 1 illustrates use of beacons in a method according to the present invention; and,

Figure 2 illustrates a frequency band structure for the beacons and managed bands using the method of the present invention.

Fig. 1 illustrates one example where the method of the present invention may be 25 applied. A first receiver 1 which is trying to receive from a first transmitter (not shown) needs to prevent interference 2 on the frequency that it is trying to receive. To do this it transmits a beacon 3 representing that frequency. This beacon is received at a beacon receiver associated with a second transmitter 4 which is trying to transmit to a second receiver (not shown). If the second transmitter 4 transmits it will produce 30 interference 2 at the first receiver 1.

The beacon 2 transmitted from the first receiver represents the frequency at which the receiver 1 is trying to receive. Clearly this cannot be the same frequency as that being received, so a nearby frequency is assigned for this purpose. A compromise

is required, so that the frequency used is near enough to have good correlation of path loss, but far enough away to allow diplexing.

Although Fig.1 only shows one receiver and one transmitter, in a practical system there are multiple receivers and transmitters. The basic principle is that some or 5 all of the receivers 1 transmit a beacon 3 at a suitable power level. Any terminal with a transmitter 4 that can hear a beacon related to a given frequency must reduce its transmitter power according to the strength of the beacon. If the permissible power is too low to allow the required communications then the terminal cannot use that frequency. A terminal operating dynamic channel allocation (DCA) would scan around 10 the beacons and select the frequency whose beacon was received at the lowest power. If multiple beacons representing the same frequency are received at different powers the power of interest is taken to be the strongest.

The beacon transmission band is arranged as a band of frequencies at either end (or both ends) of the band being managed. There is one separate beacon signal relating 15 to each frequency in the band and these must be multiplexed together. Determination of beacon power and interference limitation is explained below for the situation illustrated in Fig.1. In the two links, primary status has been allocated to a terminal trying to receive and secondary status to the one trying to transmit and for simplicity the transmitter for the primary link and the receiver for the secondary link are not 20 shown. The primary receiver 1 is transmitting its beacon 3 in order to guarantee reception.

The following assumptions are made:

- P_P is the primary receiver beacon power
25 P_S is the secondary transmitter power spectral density
 g is the path gain between the two equipment connectors – this includes propagation path gain ($<<1$), antenna gains and feeder losses
 N_P is the primary receiver noise figure
 N_S is the noise figure of the beacon receiver at the secondary transmitter
 B_S is the effective bandwidth at the beacon receiver for receiving the beacon
30 γ_S is the minimum signal to noise ratio for receiving the beacon in its effective bandwidth
 d is the factor by which we allow the primary receiver to be de-sensitised

The level of P_S needs to be set such that it is measurable at a point where a secondary transmitter could produce non-trivial interference at the primary receiver. This is done by fixing the maximum power that a secondary transmitter can emit when the beacon is undetectable, such that only acceptable interference arises.

5 The limit sensitivity for the beacon receiver is $L_S = kTB_sN_{sys}$
 (where k is Boltzmann's constant and T is the operating temperature in degrees Kelvin).

10 Thus, if the received beacon power is less than or equal to L_S , then the secondary transmitter will emit a power spectral density of P_{S-MAX} Watts/Hz. This is set so that if the beacon signal is equal to L_S , then only acceptable interference will be generated at the primary receiver. Assuming that this condition happens for a path gain of g , then the received beacon power will be given by

$$L_S = P_P \cdot g. \text{ Thus } P_P \cdot g = kTB_sN_{sys}.$$

15 The interference power spectral density generated at the primary receiver will then be $P_{S-MAX} \cdot g$. The acceptable level for this will be $kTN_p(d-1)$. Thus $P_{S-MAX} \cdot g = kTN_p(d-1)$. Eliminating g from the two equations gives

$$P_{S-MAX} = P_p(d-1) \frac{N_p}{\gamma_s B_s N_s}$$

This assumes that the operating temperature is the same at both locations which will usually be true to a reasonable approximation.

20 In practice, for universal operation, a reference needs to be set. For example, if the reference is set at 1W for a de-sensitisation of 3 dB with a 0 dB noise figure primary receiver, then conveniently this gives a beacon power for 3 dB de-sensitisation equal to $1/N_p$ Watts.

In the following example, the permitted de-sensitisation is 3 dB, both receivers
 25 have a noise figure of 6 dB, $B_s = 1$ kHz and $\gamma_I = 10$ dB. Thus, the beacon power is $\frac{1}{4}$ Watt and so $P_{S-MAX} = 2.5 \times 10^{-5}$ W/Hz.

30 In this example, a secondary transmitter with a bandwidth, for example, of 100 KHz is allowed to transmit up to 2.5 W if it cannot hear a beacon. As the beacon is detected and its received power climbs above L_S , then the maximum allowable power has to be reduced pro-rata.

From this simple analysis it can be seen that the appropriate beacon power is independent of primary system bandwidth, operating range, range to the secondary

transmitter(s) and the primary system required signal to noise ratio. The beacon power depends only on the primary receiver noise figure and the allowable de-sensitisation. Once a reference for beacon power has been legislated, only relatively small variations will result. The reference power, assuming receiver noise temperatures no lower than 5 290° K and desensitisation by no less than 3 dB, will be the maximum beacon power ever used.

Feeder loss and antenna patterns do not affect the beacon power or maximum secondary power conditions for the beacon system, although they do affect the spectrum sharing that is available. The maximum power a secondary transmitter may 10 transmit depends on the sensitivity of its beacon receiver. Thus the equipment manufacturer is motivated to produce beacon receivers with low noise figure and good 15 beacon detectability.

Fig. 2 illustrates an arrangement of a beacon band structure for the method of the present invention. A beacon signal represents the reception of a particular frequency 15 in a band. At least one beacon type is required for every frequency. It is necessary to determine how large a beacon managed band 8, 9 ought to be, how much spectrum should be allocated to beacon transmission and how many frequencies should be allocated within a beacon managed band. As in all tradeoffs the result is a compromise. On the one hand, it is desirable that the beacon managed band is as large as possible to 20 reduce the relative overhead of the beacon channel, but on the other hand it should be as narrow as possible in order to guarantee the same radio propagation characteristics across the band and, in particular, between any sub-band and the beacon transmission band.

As it is not practical to reduce the managed bandwidth down to the multipath 25 fading correlation bandwidth of the channel, then the beacon transmission band must be wide enough to have reasonable inherent multipath diversity. If this beacon reception diversity is greater than that available to the users of the band, then the interference sensing mechanism will be conservative, i.e. it will tend to overestimate the interference generated. In the alternative situation, an underestimate may occur, so it 30 may be necessary to introduce a fading margin to account for this difference. This can be done by transmitting the beacon at slightly higher power.

Based on the above requirement, it is desirable that the beacon managed band is no greater than about 5% of its centre frequency, which gives reasonable correlation of

average propagation characteristics. Some care is needed to ensure that no tight frequency selectivities in the antennas cause difficulties. Such a problem is more likely to arise in the primary service, which will probably be incumbent, than in the secondary service, since the secondary service is generally designed to use any of the frequencies
5 in the beacon managed band. One solution to the problem for the primary service is to bias the beacon transmitted power according to the relative antenna gains in the beacon transmission band and the represented signal frequency. Another feature is that the beacon transmission band is preferably set so that it does not exceed about 5% of the beacon managed band in order to avoid there being an unacceptable overhead for the
10 beacons.

Consideration must be given to the provision of duplexing filters since the primary receiver must transmit and the secondary transmitter must receive. Clearly if the beacon must be transmitted to represent a frequency that is right next to the beacon transmission band there will be a problem. This difficulty can be solved by having
15 beacon transmission bands 5, 6, 7 at each end of the beacon managed band 8, 9: The frequencies available for use are then represented by the beacon transmission band which is further away i.e. a first range 10 of frequencies in the managed band 8, are managed by the transmission band 5, whereas a second range of frequencies 11 are managed by band 6. At first this might seem to double the overhead. However, this is
20 overcome by having the beacon transmission bands manage bands on either side, hence band 6 also manages frequency range 12 of beacon managed band 9 whilst band 7 manages frequency range 13. The size of the first and second range 10, 11 is generally equal, i.e. half of each managed band is managed by one beacon transmission band 5 and half by the other 6, but these ranges can be different proportions of the total beacon
25 managed band 8, such as 1/3 and 2/3 or other ratios.

In a specific example, a frequency band is set around 2 GHz, with a bandwidth 8 of 100 MHz. This leads to a beacon transmission band 6 of 5 MHz (based on 5% bandwidth). According to the structure of Fig. 2, there is always a frequency spacing of at least 50 MHz to allow for duplexer filtering.

30 The choice of number of frequencies represented is again, a trade-off between complexity and flexibility. From the viewpoint of simplicity we would choose to have as few as possible. For example, if our 100 MHz wide band had an incumbent system with 5 MHz channels then the obvious answer would appear to be to have twenty, 5

MHz channels. However, this pre-supposes that there will be only one further tier of spectrum sharing and that the sharing system will also have a bandwidth of 5 MHz or multiples thereof. It may be that some sharing systems are introduced based on 7 MHz. In this case they may need to transmit as many as 3 beacons to keep a channel free.

- 5 This would reserve 15 MHz of spectrum which would be inefficient. From the viewpoint of flexibility the answer would be to use the highest common divisor of all anticipated channel spacings. The disadvantage of this approach is that it might lead to the need for a large number of beacons for each frequency used. For example, a spacing of 1 MHz would require 5 beacons to keep a 5 MHz channel free. The number of 10 frequencies would thus be a compromise between wasted spectrum and having a large number of beacons. At this stage a bandwidth of 5 MHz nevertheless seems reasonable leading to 20 beacons in this example.

At any given beacon receiver there may be beacons from a number of receivers, relating to some different and some common frequencies. A would-be transmitter 15 wants to examine beacons for a particular frequency. If there are many beacons for that frequency it is desirable that the task should not grow unduly. The options for multiple access are division in frequency, time and code (FDMA, TDMA and CDMA respectively).

- If beacon transmitters and receivers are to be low cost, then implicitly they have 20 relatively inaccurate frequency references. For example, a 20 ppm end to end error would not be unreasonable, but at 2 GHz this would correspond to an error of 40 kHz making any FDMA system difficult to operate. Additionally, each carrier would be relatively narrowband for FDMA. For example, representing 20 frequencies in 5 MHz gives a maximum channel spacing of 250 kHz and so the multipath diversity inherent 25 in a bandwidth of 5 MHz is not exploited.

- It is desirable for there to be no need for co-ordination between the different beacon transmitters, so no framing structures can be imposed, as required in TDMA. GPS receivers could be used to achieve this, but these can also have problems, particularly given the need for indoor operation and without some sort of timing 30 control, the only available form of multiplexing in the time domain is random transmission, akin to the Aloha protocol. However, a problem with this is that *a priori* it is not possible to be sure how many beacons need to be received in a given area, so it is difficult to dimension the average duty cycle of the transmission.

CDMA allows for a very large number of beacons relating to the same frequency since they would have the same code and would be separated only in time, but this suffers from the near far problem. If a strong beacon relating to one frequency is received then the receiver may know that that frequency is unavailable. However, it 5 cannot know whether other frequencies are available since the beacon receiver is desensitised and it must be assumed that other beacons *may* be present.

Overall, a combination of random TDMA and CDMA is preferred. The CDMA element provides the ability to measure multiple beacons relating to the same frequency and the TDMA element provides resistance to the near far effect. By listening for long 10 enough, a beacon receiver can accumulate adequate confidence that, if there had been a beacon for a particular receiver at a lower level than a strong signal then it would have had an opportunity to receive it. Where a beacon transmitter needs to send multiple signals from one location (e.g. for a frequency that requires several contiguous sub bands) they must all be transmitted simultaneously and not in sequence. This is to avoid 15 a single source occupying too much of the time.

It is desirable that the beacon receiver is made as sensitive as possible, which can be done by making the correlation period for the CDMA component as long as practically possible. The problem here is that correlating for a long period is subject to the effects of the frequency error between the beacon transmitter and the receiver. For 20 example, at 2 GHz there could be an end-to-end error of 40 kHz. The longest period over which it is usually possible to correlate is one half cycle of the error frequency, in this case 12.5 μ s. This leads to an effective bandwidth of 80 kHz which is too wide for sensitive reception. This can be radically improved by the use of an FFT correlator, for example as described in D J R Van Nee & A J R M Coenen, "New Fast GPS Code-Acquisition Technique Using FFT", IEE Electronic Letters, 17th January 1991 Vol 27, 25 pp158 – 160. By choosing the number of FFT bins, a correlation period of any length can be made. The maximum useful correlation period depends on the correlation time of the channel. This varies from about 200 μ s to 1 ms depending on the mobility of the transmitter, receiver or both. The underlying code can be set to have a duration of 30 around 200 μ s with up to, for example, 4 repetitions allowed which can be used in a static receiver to improve sensitivity. This gives an effective bandwidth from 5 kHz down to 1.25 kHz allowing good sensitivity. The preferred FFT size is 16, 32 or 64, all practical with today's technology for a 10 MHz sampling rate, although with changes

in technology other sizes may be appropriate. There is some scalloping loss from the FFT, but this can be improved by randomising the frequency over ± 0.5 bins.

A further issue for the beacon information is that although it appears that it is only necessary to know the frequency represented by the beacon, it is also necessary to 5 distinguish beacons transmitted from receivers in the same system as the transmitter, otherwise the transmitter never gets an opportunity to transmit. At any given time there are only a relatively small number of systems sharing any given band. Thus, only a small number of codes are needed to distinguish the different systems will be needed. Eight bits would normally be sufficient. Every beacon transmitter appends its system 10 type identifier to each beacon and transmits this using the applicable spread spectrum code. This part is not related to the represented frequency, but is married to it by its fine frequency, i.e. the FFT bin selected and its timing. Terminals for a given system only need to correlate codes for their own system, so a single matched filter can be used here. Thus, there is no loss in sensitivity in determining whether a beacon came from a 15 beacon receiver's own type of system. In this case, the beacon receiver ignores this particular signal when determining the power at which it could operate.

As stated earlier, for a given represented frequency the preferred power to use for the beacon is the largest of the received powers. However, in multipath, there may be several paths all relating to the same beacon transmission. Inevitably there will be 20 some ambiguity in determining this situation, but this can be resolved in most cases by only taking other multipath components from the same FFT bin output as that which identified the strongest output, which means that in the absence of significant Doppler the multipath components from any given beacon transmitter all fall on the same fine frequency. In addition, only multipath components are taken from a narrow window 25 around the largest peak, which constrains the components taken to within the maximum delay span of the propagation medium. Finally, a maximum number of paths are taken to avoid including peaks of noise.

The description above sets out the situation for a primary and a secondary system in which the receivers of the primary system have beacon transmitters and the 30 transmitters of the secondary system have beacon receivers, but this case gives no protection for the secondary system. A more general set of possibilities is shown in Table 1.

Type of System	Beacon Transmitter	Beacon Receiver
1. Primary (Dominant)	x	
2. co-operating	x	x
3. Tertiary (Unprotected)		x

Table 1

In addition, beacon transmissions could be made responsive for systems that could
 5 tolerate brief loss of communications. In this case a receiver would only transmit a beacon when it encountered unacceptable interference. The interfering sources would then detect the beacon and either stop transmitting or reduce their power. It might in some cases be desirable for type 2 systems to receive and transmit essentially simultaneously on the beacon channel. This would apply if the dynamics of spectrum
 10 sharing were such that it was not appropriate for the terminal to stop transmitting the beacon when it was making its own system transmissions. The random TDMA transmission format for the beacon would allow this to happen. The choice of level of protection for a system would be a matter for a combination of common sense and legislation.

15 There is potentially a great deal more flexibility in the actual beacon power a device may transmit. This may be influenced by a number of factors such as the manufacturer needing or wanting to increase the nominal transmitted beacon power to take account of inaccuracies in RF gains, beacon transmitters etc. Regulation would seek to prevent manufacturers degrading their equipment performance requirements
 20 excessively at the expense of spectral efficiency and for this reason there may be a case for regulating an absolute maximum beacon transmitted power for a certain class of receiver. Further, it may be desirable, and in some cases, acceptable to increase the beacon transmitted power to take account of multipath fading that sometimes reduces the received beacon power without affecting the interference path

25 There are several fundamental approaches that can be applied to the setting of the beacon power. A fixed beacon power can be transmitted, based on the equipment design. This approach is robust, but not very flexible and does not always lead to the best protection for the receiving equipment against multiple interferers. Another option

is to arrange for the primary receiver to estimate the received noise and set the beacon power accordingly. In this case, as more interferers share the frequency the interference rises. This rise reduces the margin for further increase in interference leading to an increase in the beacon power. This increase forces the interferers to

5 reduce their power, or prevents additional interferers from operating, in order to maintain the receiver sensitivity. Alternatively, the primary receiver is arranged to set its beacon power according to the maximum acceptable level of interference given the signal that it is receiving. This approach clearly should not be used in conjunction with power control within the system, as the two systems will interact in complex and

10 potentially undesirable or unpredictable ways. However, for fixed power systems this approach can be used to maximise spectrum sharing.

The fundamental beacon concept is flexible enough to allow any or all of the above approaches to apply for different systems operating in the same beacon managed band. In any given band, care needs to be taken in considering the types of system

15 sharing the spectrum to ensure that overall stability is preserved. It would be possible, for example, to operate one of each of the types of system in Table 1 without instability. The type 1 system does not respond to beacons; the type 2 system defers to the type 1 system and protects itself from the type 3 system and the type 3 system defers to the types 1 and 2 systems and uses whatever spectrum it can find. For

20 broadcast systems there are further degrees of freedom. For example, if one wished to use the beacon concept to allow sharing with the television band then every receiver needs to be equipped with a beacon transmitter, however, a beacon only needs to be transmitted corresponding to the channel that the viewer is receiving at the time.

The strong desirability for transmitting the beacon through the same antenna as

25 that used for reception, in order to share the antenna's properties, precludes the use of simple masthead pre-amplifiers. Instead, either a masthead pre-amplifier is arranged to be the source of the beacon transmission, with the reference power chosen in such a way that the absolute power requirement is always modest; or else a diplexer is built into the masthead pre-amplifier with gain for the beacon transmitter as well.

30 Although these requirements add complexity, they do not create insuperable problems and additional revenue from the freed up spectrum should more than cover such costs.

The whole concept of spectrum sharing through beacons depends on the beacons being received whenever present with power above the presumed reference

sensitivity. Equipment with a failed beacon receiver may transmit at maximum power on any frequency and cause arbitrary interference, so the status of the beacon receivers must be known at all times. To deal with this a built in self-test for the beacon receivers can be provided, in which a known beacon is injected periodically into the receiver at 5 limit range level in order to test its function. Failure to receive the test beacon results in the equipment being prevented from transmitting.